

Claims

We claim:

1. A parallel processor computer interconnect router, the interconnect router

comprising:

5 a multicasting module operable to receive a single incoming multicast packet comprising a destination identifier identifying a plurality of destination nodes, and to output multiple unicast packets, each of the multiple unicast packets comprising a destination header identifying a single destination node from among the plurality of destination nodes; and

10 a gathering module operable to receive unicast reply packets from the plurality of destination nodes, and to output a combined multicast reply packet.

2. The parallel processor computer interconnect router of claim 1, wherein the single incoming multicast packet comprises a cache invalidation message.

15

3. The parallel processor computer interconnect router of claim 2, wherein the unicast reply packets comprise cache invalidation acknowledge packets.

4. The parallel processor computer interconnect router of claim 1, wherein the output
20 combined multicast reply packet is routed to a reply destination node designated by the single incoming multicast packet.

5. The parallel processor computer interconnect router of claim 1, wherein the output combined multicast reply packet is routed to a reply destination node that is a node other than the node sending the single incoming multicast packet.
- 5 6. The parallel processor computer interconnect router of claim 1, wherein the output combined multicast reply packet is routed to the node sending the single incoming multicast packet.
7. The parallel processor computer interconnect router of claim 1, wherein the router
10 is associated with a local plurality of processors comprising a subset of processors in a parallel processor computer system, and creates multicast packets only for processors locally known to the router.
8. The parallel processor computer interconnect router of claim 1, wherein the
15 gathering module comprises a gather buffer which is allocated to gather unicast reply packets if a gather buffer is available.
9. The parallel processor computer interconnect router of claim 8, wherein the gather buffer is allocated if available on receipt of incoming multicast packets that indicate a
20 multicast with gather is desired.
10. The parallel processor computer interconnect router of claim 9, wherein incoming

multicast packets that indicate a multicast with gather is desired are converted to a multicast without gather if a gather buffer cannot be allocated.

11. A method of routing packets via a router in a parallel processing computer

5 interconnect network, comprising:

receiving in the router an incoming multicast packet comprising a destination identifier identifying a plurality of destination nodes;

outputting from the router multiple unicast packets, each of the multiple unicast packets comprising a destination header identifying a single destination node

10 from among the plurality of destination nodes; and

receiving in the router unicast reply packets from the plurality of destination nodes, and;

outputting from the router a combined multicast reply packet.

15 12. The method of claim 11, wherein the single incoming multicast packet comprises a cache invalidation message.

13. The method of claim 11, wherein the unicast reply packets comprise cache invalidation acknowledge packets.

20

14. The method of claim 11, wherein the output combined multicast reply packet is routed to a node designated by the single incoming multicast packet.

15. The method of claim 11, wherein the output combined multicast reply packet is routed to a reply destination node that is a node other than the node sending the single incoming multicast packet.
- 5 16. The method of claim 11, wherein the output combined multicast reply packet is routed to the node sending the single incoming multicast packet.
17. The method of claim 11, wherein the router is associated with a plurality of locally known processors comprising a subset of all processors in a parallel processor
- 10 computer system, and creates multicast packets only for processors locally known to the router.
18. The method of claim 11, further comprising allocating a gather buffer to gather unicast reply packets if a gather buffer is available.
- 15 19. The method of claim 18, wherein the gather buffer is allocated if available on receipt of incoming multicast packets that indicate a multicast with gather is desired.
20. The method of claim 19, further comprising converting incoming multicast
- 20 packets that indicate a multicast with gather is desired to a multicast without gather if a gather buffer cannot be allocated.

21. An information handling system comprising multiple processors connected via an interconnect network and at least one router, the router comprising:

a multicasting module operable to receive a single incoming multicast packet comprising a destination identifier identifying a plurality of destination nodes, and to
5 output multiple unicast packets, each of the multiple unicast packets comprising a destination header identifying a single destination node from among the plurality of destination nodes; and

a gathering module operable to receive unicast reply packets from the plurality of destination nodes, and to output a combined multicast reply packet.

10

22. The information handling system of claim 21, wherein the single incoming multicast packet comprises a cache invalidation message.

23. The information handling system of claim 21, wherein the unicast reply packets
15 comprise cache invalidation acknowledge packets.

24. The information handling system of claim 21, wherein the output combined multicast reply packet is routed to a node designated by the single incoming multicast packet.

20

25. The information handling system of claim 21, wherein the output combined multicast reply packet is routed to a reply destination node that is a node other than the

node sending the single incoming multicast packet.

26. The information handling system of claim 21, wherein the output combined
multicast reply packet is routed to the node sending the single incoming multicast
5 packet.

27. The information handling system of claim 21, wherein the router is associated
with a locally known plurality of processors comprising a subset of all processors in a
parallel processor computer system, and creates multicast packets only for processors
10 locally known to the router.

28. The information handling system of claim 22, wherein the gathering module
comprises a gather buffer which is allocated to gather unicast reply packets if a gather
buffer is available.

15

29. The information handling system of claim 28, wherein the gather buffer is
allocated if available on receipt of incoming multicast packets that indicate a multicast
with gather is desired.

20 30. The information handling system of claim 29, wherein incoming multicast
packets that indicate a multicast with gather is desired are converted to a multicast
without gather if a gather buffer cannot be allocated.